Bloom Filters and Compact Hash Codes for Efficient and Distributed Image Retrieval

Andrea Salvi* , Simone Ercoli, Marco Bertini and Alberto Del Bimbo
Media Integration and Communication Center, Università degli Studi di Firenze
Viale Morgagni 65 - 50134 Firenze, Italy
{*name.surname*}@unifi.it

arXiv:1605.00957v1 [cs.MM] 3 May 2016

ABSTRACT

This paper presents a novel method for efficient image retrieval, based on a simple and effective hashing of CNN features and the use of an indexing structure based on Bloom filters. These filters are used as gatekeepers for the database of image features, allowing to avoid to perform a query if the query features are not stored in the database and speeding up the query process, without affecting retrieval performance. Thanks to the limited memory requirements the system is suitable for mobile applications and distributed databases, associating each filter to a distributed portion of the database. Experimental validation has been performed on three standard image retrieval datasets, outperforming state-of-the-art hashing methods in terms of precision, while the proposed indexing method obtains a 2× speedup.

1. INTRODUCTION

Content based image retrieval (CBIR) has been an active research topic in computer vision and multimedia in the last decades, and it is still very relevant due to the emergence of social networks and the creation of web-scale image databases. Most of the works have addressed the development of effective visual features, from engineered features like SIFT and GIST to, more recently, learned features such as CNNs [2]. To obtain scalable CBIR systems features are typically compressed or hashed, to reduce their dimensionality and size. However, research on data structures that can efficiently index these descriptors has attracted less attention, and typically simple inverted files (e.g. implemented as hash tables) are used.

In this paper we address the problem of approximate nearest neighbor (ANN) image retrieval proposing a simple and effective data structure that can greatly reduce the need to perform any comparison between the descriptor of the query and those of the database, when the probability of a match is very low. Considering the proverbial problem of finding a needle in a haystack, the proposed system is able to tell when the haystack probably contains no needle and thus the search can be avoided completely.

To achieve this we propose a novel variation of an effective hashing method for CNN descriptors, and use this code to perform ANN retrieval in a database. To perform an immediate rejection of a search that should not return any result we store the hash code in a Bloom filter, i.e. a space efficient probabilistic data structure that is used to test the presence of an element in a set. To the best of our knowledge this is the first time that this data structure has been proposed for image retrieval since, natively, it has no facility to handle approximate queries. We perform extensive experimental validation on three standard datasets, showing how the proposed hashing method improves over state-of-the-art methods, and how the data structure greatly improves computational cost and makes the system suitable for application to mobile devices and distributed image databases.

2. PREVIOUS WORKS

Visual features. SIFT descriptors have been successfully used for many years to perform CBIR. Features have been aggregated using Bag-of-Visual-Words and, with improved performance, using VLAD [16] and Fisher Vectors [27].

The recent success of CNNs for image classification tasks has suggested their use also for image retrieval tasks. Babenko *et al.* [2] have proposed the use of different layers of CNNs as features, compressing them with PCA to reduce their dimensionality, and obtaining results comparable with state-of-the-art approaches based on SIFT and Fisher Vectors. Aggregation of local CNN features using VLAD has been proposed in [33], while Fisher Vectors computed on CNN features of objectness window proposals have been used in [31].

Hashing. One of the most successful visual feature hashing methods presented in the literature is Product Quantization (PQ), proposed by Jégou *et al.* [15]. In this method the feature space is decomposed into a Cartesian product of subspaces with lower dimensionality, that are quantized separately. The method has obtained state-of-the-art results on a large scale SIFT and GIST features dataset. The good performance of the Product Quantization method has led to development of several related methods that introduce variations and improvements. Norouzi and Fleet [22] have built two variations of k-means (Orthogonal k-means and Cartesian k-means) upon the idea of compositionality of the PQ approach. Ge *et al.* [9] have improved PQ minimizing quantization distortions w.r.t. space decomposition and quantization codebooks, in their OPQ method; He *et al.* [12]

*andreasalvi89@gmail.com

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
.

have approximated the Euclidean distance between codewords in k-means method, proposing an affinity-preserving technique. More recently, Kalantidis and Avrithis [17] have proposed to use a local optimization over a rotation and a space decomposition, applying a parametric solution that assumes a normal distribution, in their vector quantization method (LOPQ).

Most of recent approaches for CNN features hashing are based on simultaneous learning of image features and hash functions as in the method of Gao *et al.* [8], that uses visual and label information to learn a relative similarity graph, to reflect more precisely the relationship among training data.

Unsupervised two steps hashing of CNN features has been proposed by Lin *et al.* [21]. In the first step Stacked Restricted Boltzmann Machines learn binary embedding functions, then fine tuning is performed to retain the metric properties of the original feature space.

Indexing. Typically hashed features are stored in inverted files. A few works have studied other data structures to speed up approximate nearest neighbors. Babenko and Lempitsky [1] have proposed an efficient similarity search method that generalizes the inverted index; the method, called inverted multi-index (Multi-D-ADC), replaces vector quantization inside inverted indices with product quantization, and builds the multi-index as a multi-dimensional table. Ercoli *et al.* [7] have proposed an hashing method that improves over PQ by performing multiple assignments to k-means centroids, and have stored the hash codes in Marisa Tries to greatly compress their storage.

Bloom filter. Bloom filter and its many variants have received an extremely limited attention from the vision and multimedia community, so far. Inoue and Kise [13] have used Bloomier filters (i.e. an associative array of Bloom filters) to store PCA-SIFT features of an objects dataset more efficiently than using an hash table; they perform object recognition by counting how many features stored in the filters are associated with an object. Bloom filter has been used by Danielsson [5] as feature descriptor for matching keypoints. Similarity of descriptors is evaluated using the "union" operator. Srijan and Jawahar have proposed to use Bloom filters to store compactly the descriptors of an image, and use the filter as postings of an inverted file index in [29].

3. THE PROPOSED METHOD

In the proposed approach, differently from [8], we learn a vector quantizer separately from the CNN features, so to easily replace different and pre-trained CNN networks for feature extraction, without need of retraining. Moreover, we propose to include Bloom filters into feature indexing structures to improve the speed of queries. Bloom filters act as gatekeepers that rule out immediately, with a very limited memory cost, if a query should be completely performed or if it can be avoided. The proposed data structure is very suitable for mobile and distributed applications.

3.1 Quantization Algorithm

The proposed approach is a variation of [7], which is an efficient method for mobile visual search based on a multiple assignment k-means hashing schema (*multi-k-means*) that obtained very good results, compared to PQ, on the BIGANN dataset.

The first step of the method consists in learning a standard k-means dictionary with a small number of centroids (to maintain a low computational cost). Each centroid is associated to a bit of the hash code, that has thus length equal to the number of centroids. The bit is set to 1 if the feature is assigned to the centroid, 0 otherwise. A feature can be assigned to more than one centroid, and it is assigned to it if the distance from the centroid is less than the mean distance from all the centroids (Figure 1, top). Instead, in this work we select a fixed number N of distances and we set to 1 all the bits associated to the smaller N distances (Figure 1, bottom). In the following we refer to this method as MINx. This change has proven to be more efficient when coding CNN feature descriptors, that were used in the experiments.

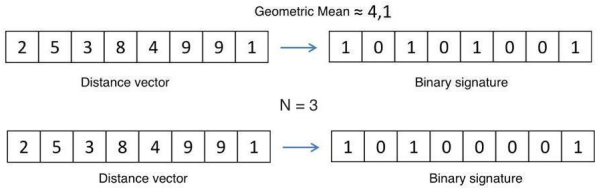

Figure 1: Binarization examples with a distance vector of 8 elements: (*top*) geometric mean (MEAN method); (*bottom*) smaller distances $N = 3$ (MINx method).

Approximate nearest neighbor retrieval of image descriptors is performed in two steps: in the first step is performed an exhaustive search over the binary codes using Hamming distances, to reduce negative effects of quantization errors. All the binary codes with Hamming distance below a threshold are selected. In the second step the candidate neighbors are ranked according to the distance computed using the full feature vector using *cosine distance*, that proved to be more effective than L_2 during the experiments.

3.2 Bloom Filter Algorithm

To improve search of feature vectors we also introduce the use of *Bloom filters* [3]. Typically this type of structures are used to speed up the answers in a key-based storage system (Figure 2).

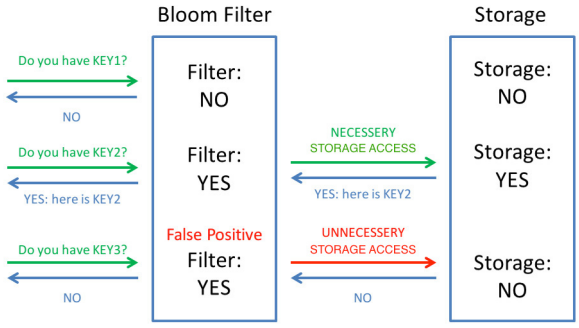

Figure 2: Memory accesses with Bloom filters

A Bloom filter is an efficient probabilistic data structure used to test if an element belongs to a set or not. This structure works with binary signatures, and can provide false positive response but not false negative and more elements are inserted into the structure and more high is the probability to obtain a false positive. To insert an element inside a

Bloom filter we need to define k hash functions which locate k positions inside the array, setting them to 1.

To check the presence of an element inside a Bloom filter we to compute the k hash functions over the element and check the related positions inside the array. If just one bit of these positions is equal to 0 it means that the element is not present inside the array; if all the checked bits are equal to 1 it means that either the element is inside the array or we have a false positive. We used the method of [19] to create the k functions from just two hash functions.

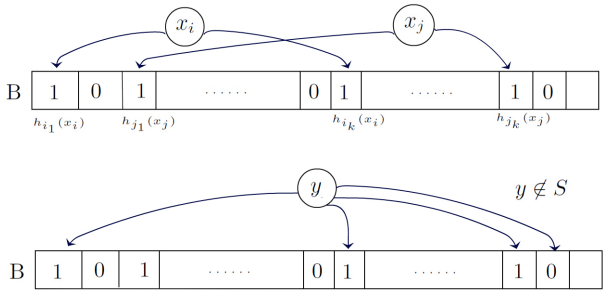

Figure 3: Example Bloom filter: (*top*) Insertion, (*bottom*) Search

A useful property of Bloom filter is that we can measure the presence of a false positive with probability:

$$(1 - e^{-kn/m})^k = (1 - p)^k = \epsilon \quad (1)$$

where m is the bit number of the array, n is the number of inserted items, p is the probability that one position of the array is equal to 0, and k is the number of hash functions. We can obtain the optimal value k which minimizes false positive probability:

$$\tilde{k} = \ln 2(m/n). \quad (2)$$

Supposed that $p = 0.5$ we can write out ϵ like

$$\epsilon = 0.5^{\tilde{k}} = (0.6185)^{m/n} \quad (3)$$

So n is strictly related to m, and in general $m = O(n)$ it is a good compromise.

Storing in the Bloom filter hash codes that are designed for ANN, as those of Sect. 3.1, results in a data structure that is similar, from a practical point of view, to distance-sensitive Bloom filters proposed in [18], where LSH functions are used as k hash functions.

3.3 Retrieval System

Our proposed retrieval system merges the methods introduced in 3.1 and 3.2. Regarding visual feature hashing we have applied the proposed method to CNNs features. Our system (Figure 4) provides a initial phase were descriptors are extracted from base images, binarized following one of the methods introduced in 3.1 and saved inside a data structure composed by a set of inverted files of hashes implementing an horizontal partition of data (allowing to distribute the database as "shard"), each one guarded by a Bloom filter. The hash code is also added to the Bloom filter of the corresponding inverted file.

During the search phase we extract the CNN descriptor from query images, compute the hash code, and check the presence of the hash in the Bloom filters, each of which guard a subset of the base. If one of this Bloom filters gives a positive response (this means that we have a positive or a false positive match), all the hash codes within an Hamming distance threshold are used to select the full feature vector. This provides a great speedup in the approximate nearest neighbor retrieval since we consider only descriptors from base coded by a Bloom filter, and below the Hamming threshold value. For each resulting original CNN descriptor we compute the distance and we rearrange results to obtain a ranked list of vectors.

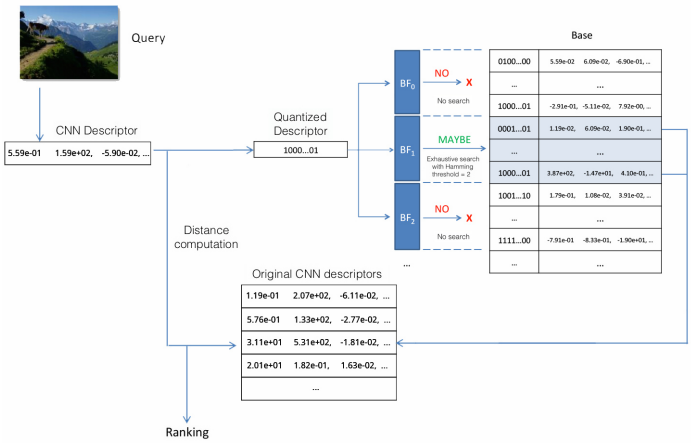

Figure 4: System overview.

4. EXPERIMENTAL RESULTS

4.1 Datasets and Configurations

We tested our system using three standard dataset: INRIA Holidays [14], Oxford 5K [23] and Paris 6K [25]. We used the query images and ground truth provided for each dataset, adding 100,000 distractor images from *Flickr 100* [24]. When testing on a dataset training is performed using the other two datasets. Features have been hashed to 64 bits binary codes, a length that has proved to be the best compromise between compactness and representativeness. Other parameters used for hashing were:

– number of N nearest distances used in the hash code computation ($N \in \{6, 10, 16, 32, 40\}$);

– Hamming distance threshold $\in \{2, 4, 6, 10, 16, 30\}$.

For the sake of brevity, in the following we report only the best combinations. For the evaluation we used the Mean Average Precision (MAP) metric. The CNN features used in the following experiments have been extracted using the 1024d average pooling layer of GoogLeNet [30], that in initial experiments has proven to be more effective than the FC7 layer of VGG [28] used in [31].

4.2 Results

In the first experiment we evaluate the effects of the method parameters, comparing the proposed hashing approach (MINx) with the original method of [7] (MEAN), a baseline that uses no hashing, and several state-of-the-art methods, among which the recent UTH method [20]. The best combinations of MINx are reported, compared on the three datasets in terms of MAP. As expected the uncompressed features perform better, but the MIN6 setup, with an Hamming distance ≥ 6 has comparable results, and greatly outperforms

any state-of-the-art hashing method. Time results in seconds, for INRIA Holidays dataset, are reported in Fig. 5. A $2\times - 10\times$ speedup can be obtained with Hamming distances between 6 and 10. Similar results, not reported here for the sake of brevity, have been obtained on Oxford 5K and Paris 6K datasets.

Table 1: MAP results on Holidays, Oxford 5K and Paris 6K datasets. The proposed MINx method outperforms all the current state-of-the-art methods. All hashes are 64 bit long.

Method	Holidays	Oxford 5K	Paris 6K
ITQ [11]	53.68	23.00	-
BPBC [10]	38.10	22.51	-
PCAHash [11]	52.80	23.90	-
LSH [6]	43.08	23.91	-
SKLSH [26]	24.09	13.39	-
SH [32]	52.22	23.24	-
SRBM [4]	51.58	21.23	-
UTH [20]	57.10	24.00	-
MIN6 Thr. 10	**75.62**	**46.03**	67.57
MIN6 Thr. 16	**75.63**	**46.06**	**67.84**
MIN10 Thr. 10	74.33	45.66	60.84
MIN10 Thr. 16	**75.63**	46.04	67.70
MEAN [7] Thr. 10	68.26	36.22	48.02
Baseline	75.63	46.06	67.84

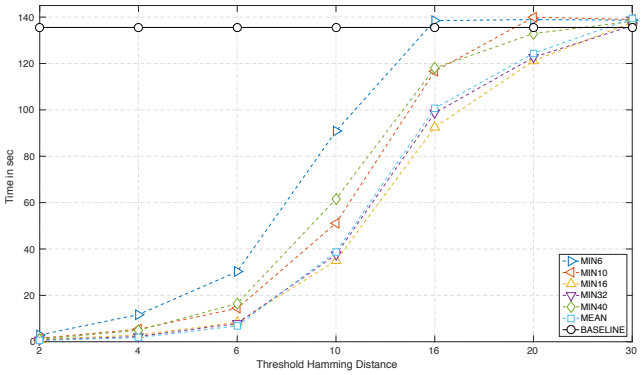

Figure 5: INRIA Holidays: time comparison for the MINx method, the MEAN [7] method and baseline without hashing.

In the second experiment we evaluate a use case in which a database of images is queried with a large number of images that do not belong to it. Hash codes have been computed with different variants of the proposed hashing method. The database contains the Paris 6K images, and it is queried with all the query images of Paris 6K and all the 100,000 distractor images. A different number of Bloom filters, with different sizes is tested and compared against a baseline that does not use any Bloom filter. MAP values and query time in seconds are reported in Tab. 2 (MAP in the first row and time in the second). The speedup obtained is about $2\times$ since a large number of distractor queries are immediately stopped by the system; the slight increase in MAP is due to the beneficial effect of elimination of some false positives of the Paris 6K images, that do not result in retrieving wrong dataset images.

Table 2: MAP+time (secs.) obtained on Paris 6K with the proposed system with different numbers of Bloom filters (1, 2 and 5) and with a baseline without filters. $2n$ and $5n$ are the size of the filters, where n is the number of stored elements. *Thr.* is the maximum Hamming distance used for hash code retrieval.

	No BF	1 BF		2 BF		5 BF	
		2n	5n	2n	5n	2n	5n
MIN6	67.57	67.57	67.57	67.94	**70.96**	67.53	68.21
Thr. 10	460.31	242.72	173.46	311.26	**205.08**	382.14	366.61
MIN10	67.69	67.69	67.69	68.37	68.46	67.66	66.70
Thr. 16	553.79	307.46	174.53	459.74	272.31	543.65	430.06

In the third experiment we evaluate a more challenging and large scale experiment: three datasets composed by distractor images and Holidays, Paris 6K and Oxford 5K images are built and stored in the proposed data structure. The standard dataset query images are then used to query the system. In this case we have used 10 filters to "shard" the database that, thus, can be distributed. Tab. 3 reports the results in terms of MAP and time (secs.). For the sake of space we report only results for MIN6 and Hamming threshold 10. Using the proposed method results in speed improvement of $2\times$ while improving MAP, except the Holidays dataset that only improves speed. The size of each Bloom filter is $\sim 6 - 62$ KB, allowing the use of the method in a mobile environment, by distributing the Bloom filters to the mobile devices and maintaining the shards of the database on the backend.

Table 3: MAP+time (secs.) obtained on Paris 6K, Oxford 5K and Holidays with the proposed system with 10 Bloom filters of varying size and with a baseline without filters. The database contains 100,000 distractor + the database images of each dataset.

# BF	Paris 6K	Oxford 5K	Holidays
No BF	58.52	42.05	**59.56**
	3.35	2.35	56.71
10 BF	59.44	41.45	52.15
10n	2.66	1.95	47.21
10 BF	61.21	42.01	42.36
20n	1.93	1.53	36.36
10 BF	62.82	**42.29**	39.26
30n	1.43	1.25	31.37
10 BF	**63.52**	**42.24**	34.29
50n	**1.21**	**1.11**	**26.59**

5. CONCLUSIONS

In this paper we have presented a simple and effective method for CNN feature hashing that outperforms current state-of-the-art methods on standard datasets. A novel indexing structure, where Bloom filters are used as gatekeepers to inverted files storing the hash codes, results in a $2\times$ speedup for ANN, without loss in MAP.

6. REFERENCES

[1] A. Babenko and V. Lempitsky. The inverted multi-index. In *Proc. of CVPR*, 2012.

[2] A. Babenko, A. Slesarev, A. Chigorin, and V. Lempitsky. Neural codes for image retrieval. In *Proc. of ECCV*, 2014.

[3] B. H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13(7):422–426, 1970.

[4] V. Chandrasekhar, J. Lin, O. Morere, A. Veillard, and H. Goh. Compact global descriptors for visual search. In *Proc. of DCC*, 2015.

[5] O. Danielsson. Category-sensitive hashing and Bloom filter based descriptors for online keypoint recognition. In *Proc. of SCIA*, 2015.

[6] M. Datar, N. Immorlica, P. Indyk, and V. S. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *Proc. of SoCG*, 2004.

[7] S. Ercoli, M. Bertini, and A. Del Bimbo. Compact hash codes and data structures for efficient mobile visual search. In *Proc. of ICME Workshops*, 2015.

[8] L. Gao, J. Song, F. Zou, D. Zhang, and J. Shao. Scalable multimedia retrieval by deep learning hashing with relative similarity learning. In *Proc. of ACM MM*, 2015.

[9] T. Ge, K. He, Q. Ke, and J. Sun. Optimized product quantization for approximate nearest neighbor search. In *Proc. of CVPR*, 2013.

[10] Y. Gong, S. Kumar, H. Rowley, and S. Lazebnik. Learning binary codes for high-dimensional data using bilinear projections. In *Proc. of CVPR*, 2013.

[11] Y. Gong and S. Lazebnik. Iterative quantization: A procrustean approach to learning binary codes. In *Proc. of CVPR*, 2011.

[12] K. He, F. Wen, and J. Sun. K-means hashing: An affinity-preserving quantization method for learning binary compact codes. In *Proc. of CVPR*, 2013.

[13] K. Inoue and K. Kise. Compressed representation of feature vectors using a Bloomier filter and its application to specific object recognition. In *Proc. of ICCV Workshops*, 2009.

[14] H. Jegou, M. Douze, and C. Schmid. Hamming embedding and weak geometric consistency for large scale image search. In *Proc. of ECCV*, 2008.

[15] H. Jégou, M. Douze, and C. Schmid. Product quantization for nearest neighbor search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(1):117–128, 2011.

[16] H. Jegou, F. Perronnin, M. Douze, J. Sanchez, P. Perez, and C. Schmid. Aggregating local descriptors into compact codes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(9):1704–1716, 2012.

[17] Y. Kalantidis and Y. Avrithis. Locally optimized product quantization for approximate nearest neighbor search. In *Proc. of CVPR*, pages 2329–2336, 2014.

[18] A. Kirsch and M. Mitzenmacher. Distance-sensitive Bloom filters. In *Proc. of ALENEX*, 2006.

[19] A. Kirsch and M. Mitzenmacher. Less hashing, same performance: Building a better Bloom filter. *Random Structures & Algorithms*, 33(2):187–218, 2008.

[20] J. Lin, O. Morère, J. Petta, V. Chandrasekhar, and A. Veillard. Tiny descriptors for image retrieval with unsupervised triplet hashing. *arXiv preprint arXiv:1511.03055*, 2015.

[21] K. Lin, H.-F. Yang, J.-H. Hsiao, and C.-S. Chen. Deep learning of binary hash codes for fast image retrieval. In *Proc. of CVPR*, 2015.

[22] M. Norouzi and D. Fleet. Cartesian k-means. In *Proc. of CVPR*, 2013.

[23] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *Proc. of CVPR*, 2007.

[24] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *Proc. of CVPR*, 2007.

[25] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. Lost in quantization: Improving particular object retrieval in large scale image databases. In *Proc. of CVPR*, 2008.

[26] M. Raginsky and S. Lazebnik. Locality-sensitive binary codes from shift-invariant kernels. In *Proc. of NIPS*, pages 1509–1517, 2009.

[27] J. Sánchez, F. Perronnin, T. Mensink, and J. Verbeek. Image classification with the Fisher vector: Theory and practice. *International Journal of Computer Vision*, 105(3):222–245, 2013.

[28] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proc. of ICLR*, 2015.

[29] K. Srijan and C. V. Jawahar. Towards exhaustive pairwise matching in large image collections. In *Proc. of ECCV*, 2012.

[30] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proc. of CVPR*, June 2015.

[31] T. Uricchio, M. Bertini, L. Seidenari, and A. Del Bimbo. Fisher encoded convolutional bag-of-windows for efficient image retrieval and social image tagging. In *Proc. of ICCV Workshops*, 2015.

[32] Y. Weiss, A. Torralba, and R. Fergus. Spectral hashing. In *Proc. of NIPS*, 2009.

[33] J. Yue-Hei Ng, F. Yang, and L. S. Davis. Exploiting local features from deep networks for image retrieval. In *Proc. of CVPR Workshops*, June 2015.